LEVEL II

AD A090247

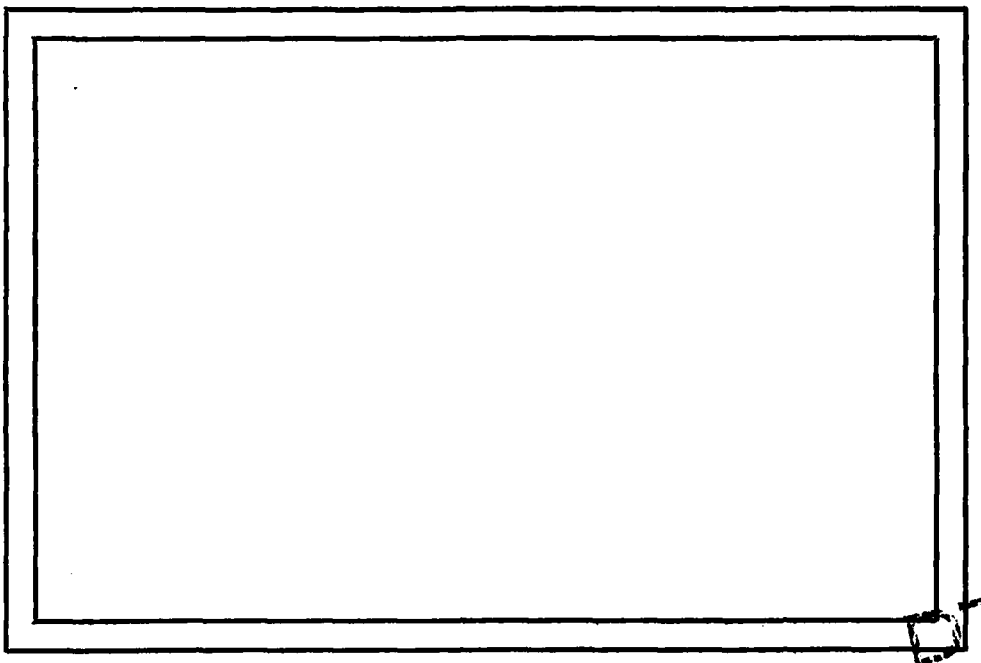

DTIC
ELECTE
OCT 1 4 1980
S
D
C

UNIVERSITY OF MARYLAND
COMPUTER SCIENCE CENTER

COLLEGE PARK, MARYLAND
20742

DDC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

80 10 3 109

(12)

(13) 141

(14) TR-894
DAAG-53-76C-0138

(11) April, 1980

(6) USING QUADTREES TO SMOOTH IMAGES

(10) Sanjay Ranade and Michael Shneier
Computer Vision Laboratory
Computer Science Center
University of Maryland
College Park, MD 20742

DTIC
ELECTE
OCT 14 1980

(15) DAAG53-76-C-0138
DARPA Order-3206

(9) Technical Rept.

ABSTRACT

Three methods for smoothing images are presented. All three use variable numbers of picture points over which the smoothing functions are defined. The first method examines the histogram of a neighborhood of each point. The subset of gray values most similar to that of the point is used in calculating the smoothing functions. The other two methods use quadtree approximations to smooth the image. One of these uses the sizes of the leaves in the quadtree to determine neighborhood sizes over which to apply a smoothing function, while the other refines the gross smoothing defined by the quadtree. All the methods perform fairly well, but the quadtree methods are particularly attractive because of the information about region sizes and homogeneity provided by the quadtree structure.

The support of the U. S. Army Night Vision Laboratory and the Defense Advanced Research Projects Agency under Contract DAAG-53-76C-0138 (DARPA Order 3206) is gratefully acknowledged.

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

411074

Accession For
GRA&I
By
Distribution/
Availability Codes
Dist Special

1. Introduction

Digital images usually contain noise of various kinds. Most image procesing tasks are considerably simplified if this noise can be removed. A common approach to noise removal is to "smooth" the image by replacing each pixel value by a new value which is a function of the values in some neighborhood of the pixel [7] .

The effect of smoothing is to make the distribution of gray levels in the image more clustered. This is reflected by the spikier histogram of the smoothed image (e.g. Figure 5) and by busyness measures which decrease when the image is smoothed. For many purposes, preprocessing an image by smoothing can significantly improve the results of later operations on the image. For example, smoothing reduces spurious edge responses to an edge operator, while in segmentation schemes based on complementing edge information with region information [3], smoothing reduces the number of false responses.

Many smoothing schemes have been proposed based on examining fixed neighborhoods of each pixel. For example, median filtering replaces each pixel value by the median of the values in some neighborhood of the pixel. Other methods attempt to vary the size of the neighborhood dynamically, for example, by adding only those points whose gray levels are sufficiently similar to the point being smoothed. In both these kinds of methods, the neighborhood size is usually kept small. In the first case, small neighborhoods preserve detail better and are less likely to cross region boundaries and blur edges. In the second case, the amount

of computation may become excessive; in addition such methods are not applicable to textured regions. It is desirable to be able to use large neighborhoods in homogeneous regions and small neighborhoods near the region boundaries. Further advantages can be gained if the cost of using large neighborhoods can be reduced. This paper describes three methods of image smoothing that exhibit these desired properties.

The first method examines a large neighborhood of a point and estimates the most suitable gray level mapping for the point by analyzing the histogram of that region. The other two methods are based on quadtree approximations of the image[5]. A $2^n \times 2^n$ gray level image can be represented in a quadtree by successively subdividing it into four quadrants or blocks. Those blocks that are reasonably homogeneous give rise to leaf nodes in the tree while those that are busy or through which edges pass are represented by branch nodes and are further subdivided. The gray level of a leaf node is the average gray level of the corresponding block in the original image. By traversal of the quadtree it is possible to reconstruct an approximation to the original image by setting the area corresponding to each leaf node to the gray level of that leaf. The size of each leaf in the tree depends on the busyness of the corresponding region of the original image. Relatively homogeneous regions are represented by large leaves while busy ones are represented by small leaves.

The second smoothing method uses the quadtree to determine the neighborhood size for smoothing. If a pixel in the image is part

of a region represented by a large leaf it is smoothed using a large neighborhood; if it is part of a region represented by a small leaf it is smoothed using a small neighborhood.

The third smoothing method uses the image reconstructed from the quadtree and a mapping function. The function maps points in the image to gray levels obtained from the quadtree approximation, and attempts to determine the best mapping for each pixel.

All three methods are discussed below and examples are given of their use.

2. Smoothing methods

2.1 Local histogram peak sharpening

For certain classes of images, peaks in the image histogram correspond to relatively large regions in the image. Usually, however, the population in a region has a small amount of variability, so that the bins of the histogram adjacent to the peaks often correspond to parts of the same region. If the adjacent values are mapped into the peak value , the histogram becomes more spike-like. This process is called peak sharpening and has been used for image segmentation [4,6,8]. More specifically, the process works as follows:

First, peaks in the histogram, h, of an image are identified. Gray value i is a peak if $h(i) > h(i+1)$ and $h(i) > h(i-1)$. Second, if h(i) is a peak, then h(i-1) and h(i+1) are mapped into h(i). This procedure can be iterated, or the definition of a peak

can be extended so that h(i) is a peak if h(i) > h(i+r) and h(i) > h(i-r), for r=1, 2, ... Eventually, only a few peaks (spikes) will remain, and can be used to segment the image. If, however, this mapping is applied to a complex image, much of the detail will be lost, and features may become distorted. This is because the mapping is determined by a global analysis of the image.

Suppose, however, that the histogram peak sharpening process is applied locally in a neighborhood of each point. That is, a histogram is constructed of a neighborhood about the point, and is sharpened as described above. The gray level mapping is then applied, mapping the point to the locally most likely gray level, without having to grow the compatible region explicitly. Figure 2 shows examples of applying the method to the images in Figure 1. While the results are fairly good and most of the detail is preserved, the method is expensive because a large neighborhood size is needed in order to construct a reliable histogram.

2.2 Smoothing using variable neighborhood sizes

Methods such as the one described above require the examination of large neighborhoods about each point, while other smoothing operations, such as those using the mean or median of a neighborhood, work best using small neighborhoods. It is to be expected that methods using variable neighborhood sizes would work better than fixed neighborhood methods. In particular, regions of an image that are reasonably homogeneous would be smoothed more effectively using a large neighborhood, while busier regions might

require smaller neighborhood sizes.

Suppose that a quadtree approximation to the image is constructed, and the distribution of leaf sizes in the quadtree is used as a guide in determining the neighborhood sizes to be used in smoothing. The size of a leaf in the quadtree corresponds directly to the degree of homogeneity of the neighborhood of the image that it represents. Large leaves correspond to large homogeneous regions, and indicate that the gray values in these regions can be smoothed using large neighborhoods. On the other hand, edges are represented by small leaves, and this indicates that a smaller neighborhood should be used. The neighborhood sizes that were used in the examples of Figure 3 to correspond to the various leaf sizes are shown in Table 1.

Note that if a median filter is used for smoothing there should be little blurring even at points corresponding to the borders of leaves (although there may be some blurring at the corners of leaves). Even an averaging function should not cause blurring because it is unlikely that a whole region will be represented by a single quadtree leaf. It is more likely that neighboring, and perhaps smaller, leaves will have similar gray levels because they represent parts of the same region, perhaps closer to the border.

This method is similar to a "region growing" method that uses the standard deviation or gradient magnitude of the gray values to limit the region sizes. The advantage of the current method is the computational efficiency provided by the quadtree structure.

2.3 Smoothing by refining the quadtree approximation

A naive smoothing method would be to construct a quadtree from an image, and then replace each pixel by the gray level of the leaf to which it corresponds. This gives the so-called 'q-image' (Figure 4a). The constant gray level across each leaf gives a blocky appearance to the image and introduces false edges.

A more satisfactory mapping can be found by taking global statistics from the image and a corresponding quadtree. Every point in the image maps into some leaf in the quadtree. The frequency with which gray level i in the image maps into gray level j in the quadtree can be used to indicate which value is to replace gray level i in the smoothed image. The most suitable mapping for gray level i is that gray level j to which i most frequently maps across the whole image (or in some neighborhood of the point). This mapping procedure reduces the strong dependence on the particular leaf to which a pixel corresponds, and this reduces the blockiness of the smoothed image. Figure 4b shows the results of applying this process to the images of Figure 1.

Note that the process can be repeated for several different quadtree approximations of the image, using different homogemeity criteria. A cumulative frequency array can be maintained, and this would reduce the effect of the somewhat arbitrary nature by which the homogeneity criteria are chosen for building the quadtree. In addition, the mapping is likely to be more valid if it is derived from several quadtree approximations.

3. Discussion

The three methods discussed above all use dynamically variable numbers of picture points to smooth an image. The first method selects points from the neighborhood by examining a histogram, while the other two methods rely on quadtree approximations of the image. The second method takes advantage of the quadtree structure to locally alter the neighborhood sizes, while the third uses a global frequency table to refine the quadtree approximation to the image.

The first of the methods, while perhaps giving the best results, is computationally expensive because of the large number of points needed to calculate the histogram. This method can be compared with local histogram flattening techniques [1,2], whose purpose is to increase contrast and separate out areas with slightly different gray levels.

The second method is more efficient because the neighborhood sizes are easily found, and only simple operations such as finding the median or average of the gray levels are required. The results are not significantly different from those of the first method, so that, on the whole, it is probably preferable to use this method.

The third method is simpler than the other two, but it occasionally misclassifies points, especially when the mapping frequency is calculated over the whole image and then applied locally at each point. It would be more expensive to compute local frequency

mappings (say on regions whose sizes are proportional to the quadtree leaf sizes) but this would undoubtedly lead to better results.

4. Conclusions

This paper has presented three methods of smoothing images using variable numbers of picture points to calculate the smoothing function. While all three methods give good results, the methods using quadtrees are attractive because of the information the quadtree provides about the homogeneity of the regions in the image and about the sizes of neighborhoods that are appropriate at different places in the image.

References

1. R. C. Gonzalez and B. A. Fites, Gray-level transformations for interactive image enhancement, presented at The Second Conference on Remotely Manned Systems, June 1975.

2. D. J. Ketcham, R. W. Lowe, and J. W. Weber, Image Enhancement Techniques for Cockpit Displays, TR-P74-530R, Display Systems Laboratory, Hughes Aircraft Co., Culver City, Calif., 1974.

3. D. Milgram, Region extraction using convergent evidence, Computer Graphics and Image Processing 11, 1979, 1-13

4. S. Peleg, Iterative histogram modification, 2, IEEE Trans. Syst., Man, Cybern. SMC-8, 1978, 555-556.

5. S. Ranade, A. Rosenfeld and J. M. S. Prewitt, Use of quad-trees for image segmentation, TR-878, Computer Science Center, University of Maryland, College Park, Maryland.

6. S. Ranade, Representation of gray level images by quadtrees, in preparation.

7. A. Rosenfeld and A. C. Kak, Digital Picture Processing, Academic Press, NY, 1976.

8. A. Rosenfeld and L. S. Davis, Iterative histogram modification, IEEE Trans. Syst. Man, Cybern. SMC-8, 1978, 300-302.

Leaf size n x n	Neighborhood Size m x m
64	13
32	11
16	9
4	5
2	3

Table 1. The neighborhood sizes corresponding to various leaf sizes.

Figure 1. Two images

Figure 2. Results of local histogram sharpening

Figure 3. Results of variable-neighborhood smoothing

Figure 4. Results of smoothing using Q-image

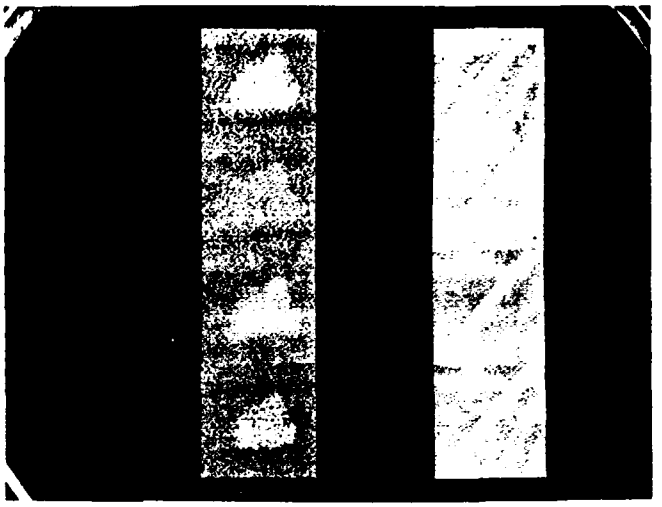

a b

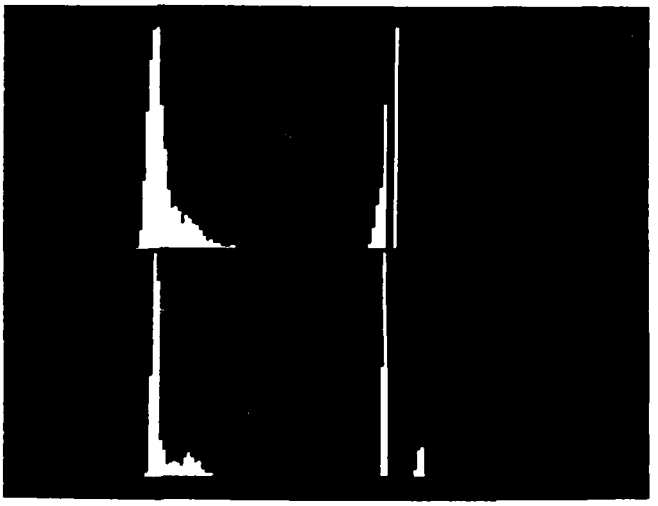

c d

Figure 5 (a-d) Histograms of tank images in Figures 1-4

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO. AD-A090 249	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) USING QUADTREES TO SMOOTH IMAGES		5. TYPE OF REPORT & PERIOD COVERED Technical
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Sanjay Ranade Michael Shneier		8. CONTRACT OR GRANT NUMBER(s) DAAG-53-76C-0138
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Vision Laboratory, Computer Science Center, University of Maryland, College Park, MD 20742		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS U. S. Army Night Vision Laboratory Ft. Belvoir, VA 22060		12. REPORT DATE April, 1980
		13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Image processing
Smoothing
Histogram sharpening
Quadtree approximation

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Three methods for smoothing images are presented. All three use variable numbers of picture points over which the smoothing functions are defined. The first method examines the histogram of a neighborhood of each point. The subset of gray values most similar to that of the point is used in calculating the smoothing functions. The other two methods use quadtree approximations to smooth the image. One of these uses the sizes of the leaves in the quadtree to determine neighborhood sizes over which to apply a smoothing

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)